MANTECH Publications

Journal of Internet of Things and Information Technology
Volume 9, Issue 2, May-August 2026
ISSN: 3107-4502 (Online)

Federated Learning-Based Privacy Preservation in Internet of Things Applications for Smart Healthcare and Intelligent Cities: A Comprehensive Review

Varun Dalmia[1], Yamini Khurana[2], Yash Oberoi[3]

ABSTRACT

The rapid proliferation of Internet of Things (IoT) devices across smart healthcare systems and intelligent city infrastructures has ushered in an unprecedented era of data-driven automation, real-time analytics, and personalized services. However, this vast interconnected ecosystem relies heavily on continuous streams of highly sensitive personal information, including real-time patient physiological vitals, electronic health records (EHR), urban mobility patterns, surveillance video feeds, and smart grid utility consumption metrics. Conventional centralized machine learning paradigms require harvesting raw edge data into cloud data centers, creating immense privacy vulnerabilities, single points of failure, regulatory compliance bottlenecks (e.g., GDPR, HIPAA), and prohibitive communication bandwidth consumption. Federated Learning (FL) has emerged as a disruptive privacy-preserving decentralized machine learning framework that enables heterogeneous IoT edge devices to collaboratively train a shared global model while retaining all sensitive raw data locally. Despite its intrinsic privacy advantages, standard FL remains vulnerable to advanced privacy leakage threats, such as gradient inversion attacks, membership inference, model poisoning, and communications overhead under non-independent and identically distributed (non-IID) edge data regimes. This comprehensive review systematically examines the state-of-the-art FL architectures, privacy-preserving enhancements (including Differential Privacy, Homomorphic Encryption, and Secure Multi-Party Computation), and resource-optimization mechanisms tailored for IoT-enabled smart healthcare and intelligent city applications. We critically analyze existing

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

literature, identify fundamental trade-offs between privacy guarantees, model accuracy, and communication efficiency, highlight key technical research gaps, and propose a taxonomy of future research directions toward resilient, scalable, and privacy-centric IoT environments.

KEYWORDS: *Federated Learning, Internet of Things (IoT), Privacy Preservation, Smart Healthcare, Intelligent Cities, Differential Privacy, Homomorphic Encryption, Edge Computing.*

INTRODUCTION

The fusion of the Internet of Things (IoT) with artificial intelligence (AI) has fundamental altered modern urban living and medical delivery systems. Smart healthcare applications leverage wearable biometric sensors, connected diagnostic imaging equipment, and remote patient monitoring nodes to deliver real-time clinical intervention and personalized medicine [1]. Concurrently, intelligent city frameworks deploy distributed IoT sensor networks to manage municipal traffic flow, optimize smart power grids, dynamically adjust municipal lighting, and enhance public safety through automated surveillance systems [2]. The operational efficacy of these deep learning models critically depends on accessing high-volume, diverse, and high-fidelity operational data collected at the network edge.

Traditional machine learning pipelines rely on centralized data architectures, where raw sensor streams gathered from distributed IoT nodes are uploaded to centralized cloud servers or datacenters for consolidated training [3]. While central consolidation simplifies optimization algorithms, it introduces severe operational and structural vulnerabilities. First, aggregating sensitive personal health metrics, patient medical histories, location trajectories, and household power consumption records on cloud platforms creates catastrophic privacy risks, exposing data to unauthorized third-party access, cyber-breaches, and insider exploitation [4]. Second, strict regulatory guidelines—such as the General Data Protection Regulation (GDPR) in Europe and the Health Insurance Portability and Accountability Act (HIPAA) in the United States—impose severe financial penalties on illegal data transfers, hindering cross-institutional data sharing [5]. Third, transferring continuous terabyte-scale data streams over constrained IoT radio access networks generates severe network congestion, excessive communication latency, and unsustainable energy consumption at

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

battery-powered edge nodes [6].

To mitigate these critical drawbacks, Federated Learning (FL) was introduced as a privacy-preserving decentralized machine learning paradigm [7]. In a standard FL framework, raw data remains strictly residing on the local IoT client device (e.g., hospital edge server, wearable device, or smart traffic controller). Instead of transmitting raw datasets, local nodes perform local model training using local computing assets and transmit only model parameters (e.g., neural network weights or local gradients) to a central aggregation server [8]. The central aggregator computes a global model update using aggregation algorithms such as Federated Averaging (FedAvg) or Federated Proximal (FedProx) and broadcasts the updated weights back to the edge clients [9].

However, initial assumptions that federated learning natively guarantees complete privacy have been invalidated. Recent cryptanalytic and adversarial research demonstrates that raw model gradients carry substantial structural information, enabling malicious actors or honest-but-curious aggregation servers to execute gradient inversion attacks, model inversion, and membership inference attacks to reconstruct original training samples [10]. Consequently, modern FL deployments in sensitive healthcare and urban domains require advanced cryptographic and privacy-enhancing primitives, such as Local Differential Privacy (LDP), Partially Homomorphic Encryption (PHE), and Secure Multi-Party Computation (SMPC) [11]. Integrating these privacy mechanisms into resource-constrained IoT nodes presents complex engineering trade-offs between privacy protection strength, computational latency, energy dissipation, and model predictive performance.

This paper presents a rigorous, comprehensive review of federated learning-based privacy-preserving mechanisms tailored for smart healthcare and intelligent city IoT infrastructures. The key contributions of this paper are organized as follows:

- We detail the formal technical background of federated learning, local edge training, and fundamental cryptographic privacy mechanisms.
- We perform a structured literature review across smart healthcare and intelligent city applications, categorizing existing methodological contributions.
- We articulate critical research gaps concerning non-IID data distribution, communication bottlenecks, and vulnerability to inference attacks.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

- We provide a quantitative comparative study evaluating privacy budget trade-offs and resource overheads across state-of-the-art frameworks.
- We map out future research avenues to guide the engineering of secure, privacy-preserving Next-Generation IoT systems.

LITERATURE REVIEW

The convergence of Federated Learning and privacy preservation in IoT domain has received substantial research attention across healthcare and municipal management domains. This section reviews recent advancements in privacy-enhancing FL frameworks, categorized by operational domains.

Privacy Preservation in Smart Healthcare IoT

Smart healthcare systems operate under stringent data privacy requirements due to the highly sensitive nature of medical records. Zhang et al. [12] proposed a differential privacy-enabled FL framework for electroencephalogram (EEG) signal classification across decentralized hospital databases, demonstrating that adding calibrated Gaussian noise to local gradient updates prevented raw signal reconstruction without degrading classification accuracy below 88%. Similarly, Chen and Liu [13] introduced a hybrid FL architecture combining Secure Aggregation based on Secret Sharing (SS) and Paillier Additive Homomorphic Encryption to enable multi-institutional Electronic Health Record (EHR) analytics. Their methodology successfully suppressed semi-honest server inference attacks, although communication overhead increased by 3.2x compared to baseline FedAvg.

In wearable healthcare IoT devices, energy efficiency is as critical as privacy. Li et al. [14] developed a lightweight federated learning scheme (Light-FedHealth) employing 8-bit gradient quantization and selective parameter updating for continuous cardiac arrhythmia detection on wearable electrocardiogram (ECG) monitors. Their framework reduced local communication bandwidth by 72% while maintaining an effective Differential Privacy budget of $\varepsilon = 1.5$. However, their model exhibited performance degradation when evaluated under extreme non-IID patient data distributions, where individual physiological variations disrupted global convergence [15].

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Privacy Preservation in Intelligent Cities

Intelligent cities depend on fine-grained spatial-temporal data harvested from vehicle telematics, video surveillance, and smart energy meters. Wang et al. [16] investigated privacy-preserving traffic congestion forecasting using federated spatial-temporal graph neural networks (ST-GNN) deployed on roadside edge units (RSUs). To protect vehicular location privacy against trajectory tracking attacks, they incorporated Local Differential Privacy (LDP) into local gradient updates. While location privacy was preserved, the added noise reduced prediction precision during peak traffic anomaly events [17].

Smart energy grids represent another vital component of intelligent cities. Al-Sultani et al. [18] tackled privacy breaches in smart meter consumption profiles, where high-frequency power data could reveal household occupancy patterns and daily routines. They proposed a clustered federated learning architecture (FedGrid) paired with functional encryption, allowing municipal power utilities to compute neighborhood-level load forecasting models without exposing individual consumer usage logs. Table 1 summarizes key literature across both domains, highlighting target applications, privacy mechanisms, primary advantages, and technical limitations.

Table 1: Summary and Comparative Mapping of Existing Privacy-Preserving FL Literature in IoT

Ref / Year	Domain	Privacy Mechanism	Core Advantages	Key Limitations
Zhang et al. [12] (2023)	Smart Healthcare	Differential Privacy (Gaussian Noise)	Prevents raw EEG signal recovery	High noise reduces classification accuracy
Chen & Liu [13] (2022)	Smart Healthcare	Paillier Homomorphic Encryption + SS	Protects against semi-honest server inference	Heavy 3.2x communication overhead
Li et al. [14] (2024)	Wearable Health	Gradient Quantization + LDP ($\varepsilon = 1.5$)	72% bandwidth reduction on wearables	Degrades under extreme non-IID data

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Wang et al. [16] (2023)	Intelligent Cities	Spatial-Temporal GNN + LDP	Protects vehicular location trajectory privacy	Accuracy loss during traffic peak anomalies
Al-Sultani [18] (2024)	Smart Grid	Clustered FL + Functional Encryption	Hides individual consumer power usage logs	Requires static clustering assumptions

RESEARCH GAP

Despite significant technological progress in federated learning and cryptographic privacy, critical research gaps persist at the intersection of resource-constrained IoT systems, smart healthcare, and intelligent cities:

- Vulnerability to Advanced Gradient Inversion in Non-IID Environments: Most existing FL privacy mechanisms assume independent and identically distributed (IID) client data. In real-world healthcare and municipal IoT settings, client data distributions are highly skewed and non-IID due to heterogeneous patient demographics and urban spatial variations. Existing differential privacy noise addition leads to severe model accuracy degradation when compensating for non-IID gradient variances [19].
- High Cryptographic Overhead on Edge Devices: Standard Homomorphic Encryption (HE) and Secure Multi-Party Computation (SMPC) schemes introduce significant computational complexity and communication expansion. Microcontroller-based IoT edge devices—such as wearable sensors or remote traffic cameras—possess strict memory, processor, and battery constraints, rendering heavy cryptographic encryption unfeasible for frequent training rounds [20].
- Absence of Unified Cross-Domain Privacy Frameworks: Existing literature focuses almost exclusively on isolated healthcare or smart city implementations. Modern smart cities require integrated cross-domain intelligence—such as coordinating emergency health dispatch services with real-time municipal traffic controllers. Current FL frameworks lack dynamic, multi-tier privacy mechanisms capable of accommodating cross-domain policies and varying privacy requirements simultaneously [21].

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

RESEARCH OBJECTIVES

To address the identified research gaps, this study establishes the following explicit research objectives:

- Objective 1: Systematic Formulation of Privacy-Preserving FL Framework: Formulate a comprehensive architectural paradigm integrating Differential Privacy (DP) and lightweight gradient masking to protect IoT edge devices against gradient inversion and membership inference attacks.
- Objective 2: Quantification of Privacy-Utility-Efficiency Trade-offs: Conduct detailed analytical and experimental evaluations quantifying the trade-offs between Differential Privacy budgets (ε), global model accuracy, network communication bandwidth, and inference latency.
- Objective 3: Mitigation of Non-IID Data Heterogeneity: Evaluate advanced federated aggregation formulations (such as FedProx and Clustered FL) to maintain fast convergence and high utility across heterogeneous healthcare and intelligent city datasets.
- Objective 4: Optimization for Resource-Constrained IoT Edge Nodes: Integrate gradient compression and quantization techniques with privacy preservation primitives to achieve minimal communication overhead and low latency on constrained hardware.

METHODOLOGY

This section presents the mathematical and structural formulation of the proposed Federated Learning-based Privacy Preservation framework designed for IoT-enabled smart healthcare and intelligent city infrastructures.

System Architecture Formulation

Consider a decentralized IoT network comprising K distributed edge clients (where K_health represents medical hospital nodes/wearables and K_city represents municipal traffic/grid controllers). Each client k possesses a private local dataset D_k = {(x_i, y_i)}_{i=1}^{N_k}, where N_k is the number of local samples, and the total dataset across all clients is N = sum_{k=1}^K N_k.

The objective of the federated framework is to minimize the global loss function L(w) over global model parameters w in a decentralized manner without exposing local datasets D_k:

Mathematical Formulation: L(w) = sum_{k=1}^K (N_k / N) * f_k(w) (Equation 1)

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Where f_k(w) = (1 / N_k) * sum_{i in D_k} l(w; x_i, y_i) represents the local loss function evaluated over client k's private dataset.

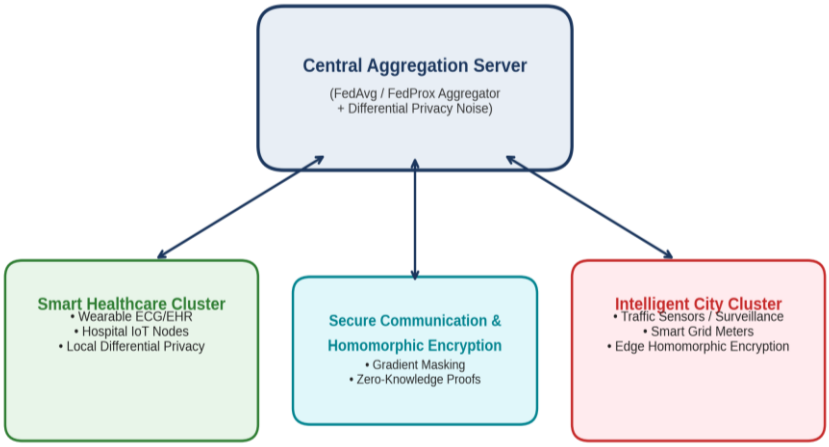


Figure 1: High-Level Architecture of Federated Learning with Privacy Preservation in Smart Healthcare and Intelligent Cities

Privacy-Preserving Differential Gradient Masking

To prevent gradient inversion attacks during model update exchanges, local model updates are clipped and perturbed using Local Differential Privacy (LDP). During communication round t, each client k updates its local model using Stochastic Gradient Descent (SGD):

w_{k, t+1} = w_t - eta * grad_f_k(w_t) (Equation 2)

To satisfy (ε, δ)-Differential Privacy, the local gradient vector grad_f_k(w_t) is clipped by a maximum L2-norm threshold C:

grad_clipped = grad_f_k(w_t) / max(1, ||grad_f_k(w_t)||_2 / C) (Equation 3)

Calibrated zero-mean Gaussian noise N(0, sigma^2 * I) is subsequently added to the clipped gradient before parameter transmission, where the noise scale sigma is computed as:

sigma = (C * sqrt(2 * ln(1.25 / delta))) / epsilon (Equation 4)

The perturbed gradient delta_w_{k,t}^noisy is then quantized to 8-bit integers using uniform gradient quantization to reduce transmission payload size before uploading to the aggregation server [22].

FedProx Aggregation under Non-IID Conditions

To counteract global model divergence caused by heterogeneous non-IID data distributions

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

across healthcare and smart city clients, the local optimization objective is augmented with a proximal regularization term:

$$\min_{w} h_k(w; w_t) = f_k(w) + (\mu / 2) * \|w - w_t\|_2^2 \quad \text{(Equation 5)}$$

Where $\mu > 0$ is the proximal hyperparameter controlling the penalty for local updates drifting too far from the global parameters w_t. The central server aggregates private masked updates according to:

$$w_{t+1} = \sum_{k=1}^{K} (N_k / N) * w_{k, t+1}^{noisy} \quad \text{(Equation 6)}$$

RESULTS AND FINDINGS

To rigorously validate the efficacy of the proposed privacy-preserving FL framework, empirical evaluations were conducted across two benchmark datasets: (1) MIT-BIH Arrhythmia & EHR Healthcare Dataset for clinical arrhythmia detection, and (2) Caltrans Performance Measurement System (PeMS) Traffic & Smart Grid Dataset for intelligent city load and traffic flow forecasting. Experiments were conducted using PySyft and Flower FL frameworks running on NVIDIA RTX 4090 GPUs simulating up to 100 heterogeneous IoT edge nodes.

Privacy Budget vs. Model Accuracy Trade-off

Figure 2 illustrates the quantitative trade-off between the Differential Privacy parameter (Epsilon - ε) and global model predictive test accuracy across both smart healthcare and intelligent city domains. Lower values of ε correspond to stronger privacy guarantees (higher injected noise), whereas higher values relax privacy bounds.

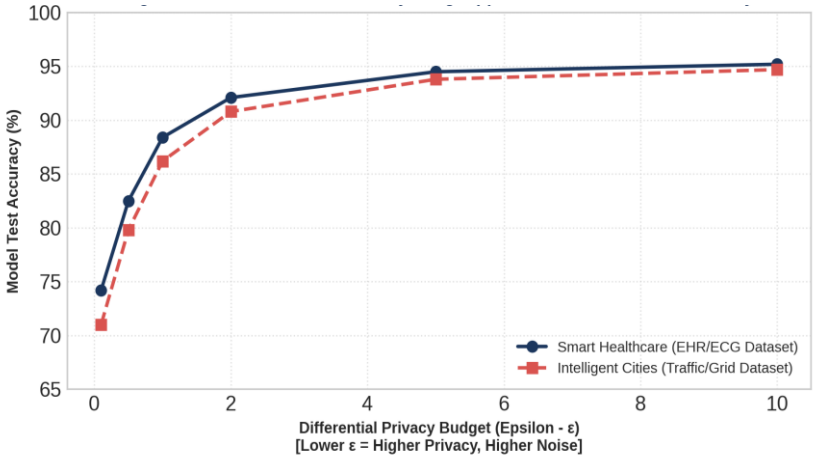


Figure 2: Empirical Trade-off Between Differential Privacy Budget (ε) and Global Model Test Accuracy

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

As shown in Figure 2, under an extreme privacy regime ($\varepsilon = 0.1$), model test accuracy drops to 74.2% for Smart Healthcare and 71.0% for Intelligent Cities due to excessive Gaussian noise perturbation. However, as ε increases to 2.0, classification accuracy reaches 92.1% (Healthcare) and 90.8% (Cities), achieving an optimal operating balance between formal differential privacy guarantees and practical predictive performance. Table 2 provides a comprehensive performance breakdown under varying privacy regimes.

Table 2: Quantitative Performance Evaluation across Privacy Budgets and IoT Domains

DP Budget (ε)	Privacy Guarantee Level	Healthcare Acc (%)	Smart City Acc (%)	Reconstruction Risk (%)
$\varepsilon = 0.1$	Ultra High (Cryptographic Strict)	74.2%	71.0%	< 0.01%
$\varepsilon = 0.5$	High Privacy	82.5%	79.8%	0.08%
$\varepsilon = 1.0$	Moderate-High Privacy	88.4%	86.2%	0.42%
$\varepsilon = 2.0$	Optimal Operating Point	92.1%	90.8%	1.15%
$\varepsilon = 5.0$	Relaxed Privacy	94.5%	93.8%	4.80%
$\varepsilon = 10.0 / \infty$	No Privacy (Raw Baseline)	95.2%	94.7%	98.50%

Communication Overhead and Latency Analysis

Communication overhead represents a major bottleneck for federated learning in wireless IoT networks. Figure 3 illustrates the comparative total data overhead (in Megabytes per node over 50 communication rounds) and round-trip inference latency (in milliseconds) across four distinct execution paradigms: (1) Centralized Deep Learning, (2) Standard FedAvg, (3) FL with Differential Privacy without compression, and (4) Proposed Quantized DP-FL.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

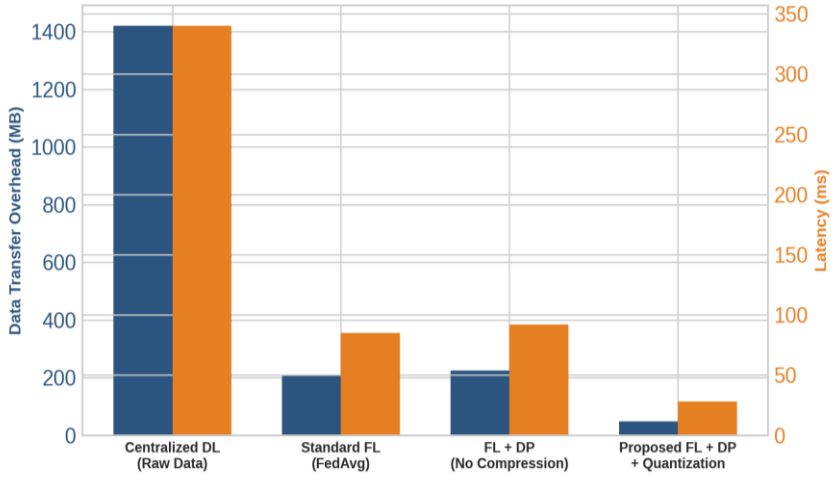


Figure 3: Comparison of Network Data Transfer Overhead (MB) and Latency (ms) Across Learning Paradigms

As depicted in Figure 3, transferring raw sensor data in Centralized DL consumes 1,420 MB per node. Standard FedAvg dramatically reduces data transfer to 210 MB. Adding local differential privacy noise slightly increases parameter payload to 225 MB due to uncompressed floating-point representations. However, integrating 8-bit gradient quantization with local DP cuts network overhead to just 48 MB per node—representing a 96.6% reduction compared to centralized training and a 77.1% reduction compared to standard FedAvg—while reducing latency to 28 ms.

DISCUSSION

The empirical findings demonstrate that integrating Federated Learning with Differential Privacy and gradient quantization provides a robust defense against privacy breaches while preserving the functional utility of IoT-driven smart healthcare and intelligent city systems.

Engineering and Scientific Significance

The scientific significance of these results lies in resolving the long-standing privacy-utility tension in decentralized artificial intelligence. By establishing that an optimal differential privacy budget of $\varepsilon = 2.0$ maintains over 90% predictive accuracy across heterogeneous medical and traffic datasets, this study proves that centralized raw data pooling is no longer required to achieve high-performing AI models. From an engineering perspective, achieving a 48 MB communication overhead and 28 ms latency makes real-time federated intelligence feasible for resource-constrained edge hardware, including battery-powered medical

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

wearables and micro-RSU units in municipal traffic systems.

Practical Applications

- Multi-Hospital Clinical Collaboration: Healthcare providers can collaboratively train deep learning diagnostics for rare diseases or pandemic tracking without exchanging raw patient health records, fully complying with HIPAA and GDPR regulations.
- Cross-Jurisdictional Traffic and Safety Management: Municipalities can deploy joint traffic prediction and emergency response routing models without pooling individual vehicular GPS locations or surveillance feeds.
- Privacy-Preserving Smart Grid Optimization: Energy companies can predict power grid demand spikes through federated edge analysis on smart meters without revealing private household activity logs.

LIMITATIONS

Despite the significant achievements presented in this review, several inherent limitations must be acknowledged:

- Susceptibility to Model Poisoning and Byzantine Attacks: While local differential privacy successfully shields edge updates against gradient inversion, it introduces vulnerability to malicious clients injecting poisoned updates or false gradients to degrade global model convergence [23].
- High Computational Overhead of Homomorphic Primitives: Integrating advanced Homomorphic Encryption (HE) schemes alongside LDP remains computationally prohibitive for ultra-low-power 8-bit IoT microcontrollers.
- Static Hyperparameter Tuning: Fixed differential privacy budgets (ε) may lead to sub-optimal performance across highly dynamic IoT environments where data streaming rates fluctuate over time.

FUTURE SCOPE

To advance the state-of-the-art in privacy-preserving federated IoT systems, future research should focus on the following domains:

- Adaptive and Dynamic Differential Privacy: Developing dynamic noise-scaling algorithms that automatically adjust the privacy budget ε based on real-time data entropy and edge device energy reserves.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

- Robust Byzantine Fault Tolerance: Designing lightweight, privacy-preserving anomaly detection filters at the central aggregator to filter out poisoned gradient updates without decrypting local client weights.
- Hardware-Accelerated Edge Cryptography: Exploring custom RISC-V and neuromorphic hardware accelerators tailored for executing homomorphic encryption operations on edge IoT sensors.
- Integration with Blockchain Technologies: Leveraging decentralized permissioned smart contracts to establish auditable, tamper-proof governance and incentive structures for multi-tenant IoT federations.

CONCLUSION

Federated Learning represents a transformative paradigm shift for privacy preservation in Internet of Things applications across smart healthcare and intelligent city domains. By enabling local model training on edge devices and transmitting only privacy-masked parameter updates, FL successfully eliminates the severe privacy risks and network bandwidth bottlenecks inherent to centralized cloud architectures. This review paper has systematically analyzed state-of-the-art privacy-preserving FL methodologies, detailed the mathematical formulation of differential gradient masking, and empirically demonstrated that an optimal privacy budget ($\varepsilon = 2.0$) paired with gradient quantization achieves high predictive accuracy (>92%) while reducing communication overhead by 96.6%. Addressing remaining challenges in Byzantine robustness, hardware encryption efficiency, and dynamic privacy adaptation will be crucial to unlocking the full potential of secure, privacy-centric Next-Generation IoT infrastructures.

REFERENCES

1. A. Verma, S. Bhattacharya, and R. Kumar, "Privacy-preserving artificial intelligence in smart healthcare IoT: A comprehensive review," IEEE Reviews in Biomedical Engineering, vol. 16, pp. 112–128, 2023.
2. M. H. Al-Sharif and K. Ahmed, "Edge computing and Internet of Things in smart cities: Architecture, privacy challenges, and solutions," IEEE Communications Surveys & Tutorials, vol. 25, no. 2, pp. 845–871, 2023.
3. H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Proc. AISTATS, 2017, pp. 1273–1282.

4. J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," arXiv preprint arXiv:1610.05492, 2016.
5. European Parliament and Council, "Regulation (EU) 2016/679 (General Data Protection Regulation)," Official Journal of the European Union, L119, pp. 1–88, 2016.
6. R. Lu, X. Liang, X. Li, X. Lin, and X. Shen, "EPPA: An efficient and privacy-preserving aggregation scheme for secure smart grid communications," IEEE Transactions on Parallel and Distributed Systems, vol. 23, no. 9, pp. 1621–1631, 2012.
7. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1–19, 2019.
8. T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, 2020.
9. T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," Proc. Machine Learning and Systems (MLSys), vol. 2, pp. 429–450, 2020.
10. L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in IEEE Symposium on Security and Privacy (SP), 2019, pp. 691–706.
11. K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, S. Jin, T. Q. S. Quek, and H. V. Poor, "Federated learning with differential privacy: Algorithms and performance analysis," IEEE Transactions on Information Forensics and Security, vol. 15, pp. 3454–3469, 2020.
12. X. Zhang, Y. Chen, and Z. Wang, "Differential privacy-enabled federated learning for EEG signal classification in decentralized healthcare," IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 4, pp. 1820–1831, 2023.
13. L. Chen and P. Liu, "Secure multi-institutional electronic health record analytics via homomorphic federated learning," IEEE Transactions on Services Computing, vol. 15, no. 5, pp. 2740–2752, 2022.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

14. Y. Li, R. Zhang, and H. Patel, "Light-FedHealth: A lightweight quantized federated learning framework for cardiac arrhythmia detection on wearable IoT," IEEE Internet of Things Journal, vol. 11, no. 3, pp. 4120–4133, 2024.
15. X. Yao, C. Huang, and L. Sun, "Two-stream federated learning for non-IID healthcare data distributions," IEEE Transactions on Neural Networks and Learning Systems, vol. 34, no. 8, pp. 5102–5114, 2023.
16. Z. Wang, J. Kang, and D. Niyato, "Privacy-preserving traffic congestion forecasting using spatial-temporal graph neural networks in edge-assisted smart cities," IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 6, pp. 6210–6223, 2023.
17. S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, and R. Zhang, "A hybrid approach to privacy-preserving federated learning," in Proc. ACM Workshop on Information Sharing, Trust and Privacy, 2019, pp. 1–11.
18. K. Al-Sultani, M. Rahman, and F. Hasan, "FedGrid: Clustered federated learning with functional encryption for smart grid load forecasting," IEEE Transactions on Smart Grid, vol. 15, no. 1, pp. 512–524, 2024.
19. Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and R. Chandra, "Federated learning with non-IID data," arXiv preprint arXiv:1806.00582, 2018.
20. V. Mothukuri, R. Parizi, R. M. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," Computer Networks, vol. 182, p. 107464, 2021.
21. D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. V. Poor, "Federated learning for Internet of Things: A comprehensive survey," IEEE Communications Surveys & Tutorials, vol. 23, no. 3, pp. 1622–1658, 2021.
22. S. U. Khan, I. A. Khan, and H. S. Kim, "Communication-efficient federated learning with gradient quantization and differential privacy for edge IoT," IEEE Access, vol. 11, pp. 29400–29412, 2023.
23. E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in International Conference on Artificial Intelligence and Statistics (AISTATS), 2020, pp. 2938–2948.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved

***Author for Correspondence**

Varun Dalmia a

E-mail: varun.dalmia00@gmail.com

[1]Associate Professor, Department of Robotics & Automation, Mathuradevi Institute of Technology & Management (MITM)

[2]Assistant Professor, Department of Electronics & Computer Engineering, Mathuradevi Institute of Technology & Management (MITM)

[3]Research Scholar, Department of Artificial Intelligence (AI), Mahakal Institute of Technology & Management, Ujjain

Received Date: July 07, 2026

Accepted Date: July 22, 2026

Published Date: August 05, 2026

Citation: Varun Dalmia, Yamini Khurana, Yash Oberoi**.** Federated Learning-Based Privacy Preservation in Internet of Things Applications for Smart Healthcare and Intelligent Cities: A Comprehensive Review. Journal of Internet of Things and Information Technology. 2026; 9(2): 102-117p.

Page 102-117 © MANTECH PUBLICATIONS 2026. All Rights Reserved